

An algorithmic approximation of the infimum reachability probability for Probabilistic Finite Automata

Sergio Giro*

*FaMAF – Universidad Nacional de Córdoba – Argentina
FCEIA – Universidad Nacional de Rosario*

Abstract

Given a Probabilistic Finite Automata (PFA), a set of states \mathcal{S} , and an error threshold $\epsilon > 0$, our algorithm approximates the infimum probability (quantifying over all infinite words) that the automata reaches \mathcal{S} . Our result contrasts with the known result that the approximation problem is undecidable if we consider the supremum instead of the infimum. Since we study the probability of reaching a set of states, instead of the probability of ending in an accepting state, our work is more related to model checking than to formal languages.

Keywords: probabilistic finite automata, reachability, automatic verification

1. Introduction

Suppose you want to analyse a system \mathcal{A} whose number of states is finite. This system reacts to inputs from the environment in a probabilistic fashion: if \mathcal{A} is in state s and receives α from the environment, the probability that \mathcal{A} transitions to state s' is $p_{s,\alpha}(s')$. Moreover, assume that the environment cannot observe the state of \mathcal{A} in order to choose the particular input α . The analysis you want to perform on this system is to calculate a tight lower bound of the probability that the system achieves a certain goal, no matter what the inputs are. For instance, inputs can model notifications of the (un)availability of resources, and you might want to check that your system sends a message with probability at least 0.8, no matter what the available resources are.

The problem in the paragraph above can be modelled using Probabilistic Finite Automata (PFA) [1, 2]. The assumption that inputs do not depend on the internal state of the state of the input is central to assert that a PFA model adequately reflects the behaviour of the system. In case the environment can observe the state of \mathcal{A} to choose the particular input α , the problem can be modelled using Markov Decision Processes (MDP) [3].

The usual semantics for PFA rely on the concept of *acceptance*, by considering the set of finite words ending in an acceptance state with probability greater than

a given cut-point η . In contrast, we focus on the concept of *reachability*, and we are interested on the probability with which each infinite word reaches some of the states in a given set \mathcal{S} . In the realm of MDPs, both the supremum and the infimum probability can be calculated in polynomial time [4]. In contrast, in the PFA setting the supremum problem is undecidable [2] for both finite and infinite words¹. In fact, the supremum probability that \mathcal{A} reaches a state in \mathcal{S} cannot be even approximated algorithmically. This undecidability result was the key to prove undecidability results for MDPs under partial information [5] as well as undecidability for Probabilistic Büchi Automata [6].

We present an algorithm to approximate the infimum probability that a PFA \mathcal{A} reaches a set of states \mathcal{S} . Moreover, the computed value v is a lower bound of the infimum and, by performing a sufficient number of iterations, we can ensure that it is as close to the infimum as desired. Using the value v , we can answer our motivating problem by stating that “the probability that the goal is achieved is at least v , no matter what the inputs are”. The fact that the value v is close to the infimum implies that the bound we provide is tight.

¹Here, we consider only infinite words, as the infimum probability over finite words is either 1, if the initial state of the system is in \mathcal{S} , or 0, if it is not.

*Urquiza 1949 16/F. Rosario, (2000) Rosario, Argentina

2. Algorithm

For our algorithm, we use the following definitions: a Probabilistic Finite Automata (PFA) is a quintuple $\mathcal{A} = (S, \Sigma, \mathcal{P}, s^i, S)$, where S is a finite set of states, Σ is a set of symbols, \mathcal{P} is a set of probability distributions on S , comprising one probability distribution $p_{s,\alpha}(\cdot)$ for each pair (s, α) in $S \times \Sigma$. The state s^i is called the *initial state* of \mathcal{A} , and S is a set of *hitting states*. We assume $s^i \notin S$.

A finite path in \mathcal{A} is a sequence

$$\pi = s^i . \alpha_1 . s_1 . \dots . \alpha_n . s_n$$

where $\alpha_i \in \Sigma$ and $s_i \in S$ for all i . Note that paths always start with the initial state s^i . We write $\text{len}(\pi)$ for n and $\text{last}(\pi)$ for s_n . In an analogous way to finite paths, infinite paths are infinite sequences alternating symbols and states. The set of all infinite paths having the finite path π as prefix is denoted by π^\uparrow .

Given a word ψ over Σ , let $\psi[k]$ denote the k -th symbol in ψ . For every infinite word ψ over Σ , for every finite path π , the probability $\Pr^\psi(\pi^\uparrow)$ is defined as 1 if $\pi = s^i$; if $\psi[\text{len}(\pi) + 1] = \alpha$, we have $\Pr^\psi(\pi . \alpha . s^\uparrow) = \Pr^\psi(\pi^\uparrow) \cdot p_{\text{last}(\pi), \alpha}(s)$; if $\psi[\text{len}(\pi) + 1] \neq \alpha$, then $\Pr^\psi(\pi . \alpha . s^\uparrow) = 0$. In the same way as for Markov chains and MDPs (namely, by resorting to the Carathéodory extension theorem), the previous definition for sets of the form π^\uparrow can be extended in such a way that, for all infinite words ψ , the value $\Pr^\psi(\mathcal{Z})$ is defined for all measurable sets \mathcal{Z} of infinite paths.

Let \mathcal{H} be the set of all infinite paths ρ such that some of the states in ρ is in S . The amount we want to approximate is $I = \inf_\psi \Pr^\psi(\mathcal{H})$. Note that \mathcal{H} can be written as

$$\mathcal{H} = \bigcup_{\pi \in C} \pi^\uparrow, \quad (1)$$

where C is the set of all finite paths π such that $\text{last}(\pi)$ is the only state of π in S .

In order to approximate I , our algorithm iterates producing two values in each iteration r . One of the values is a lower bound l_r and the other one is an upper bound u_r . These bounds comply with:

$$l_r \leq l_{r+1} \quad (2)$$

$$l_r \leq I \quad (3)$$

$$\lim_{r \rightarrow \infty} l_r = I \quad (4)$$

$$u_r \geq u_{r+1} \quad (5)$$

$$u_r \geq I \quad (6)$$

$$\lim_{r \rightarrow \infty} u_r = I. \quad (7)$$

To approximate I with error at most ϵ , the algorithm stops when $u_r - l_r < \epsilon$ (this is guaranteed to occur as both

u_r and l_r converge to the same limit), and then returns l_r . Note that u_r is also a value with error less than ϵ but, in order to give a safe lower bound on the probability that a hitting state is reached, we use the pessimistic value $l_r \leq I$.

In the next subsections, we show how to calculate upper and lower bounds complying with the desired properties.

2.1. Lower bounds

Let $\mathcal{H}_r = \bigcup_{C_r} \pi^\uparrow$ where C_r is the set of paths such that $\text{last}(\pi)$ is the only state of π in S and $\text{len}(\pi) \leq r$. By making the same observation as for Eq. 1, we deduce that \mathcal{H}_r is the set of all infinite paths reaching S after at most r symbols. We often profit from the inclusion

$$\mathcal{H}_r \subseteq \mathcal{H}_{r+1}.$$

We take $l_r = \inf_\psi \Pr^\psi(\mathcal{H}_r)$. Next, we show that this number can be calculated by brute force.

Since only the first r symbols are relevant, we need to consider each of the finite words w having exactly r symbols. The truncation operator $\psi \downarrow r$, that returns the prefix of ψ having length r , will thus be quite useful in this subsection. In addition, we use the notation $\Pr^w(\mathcal{H}_r)$ to mean $\Pr^\psi(\mathcal{H}_r)$, where ψ is any infinite word such that $\psi \downarrow r = w$.

For each w with $\text{len}(w) = r$, we construct a finite Markov chain \mathcal{M} . The procedure resembles the standard unfolding of a probabilistic automaton (or an MDP) for a particular adversary [7], and so we merely outline it. The states of \mathcal{M} are pairs (s, k) with s in S and $0 \leq k \leq r$. To describe \mathcal{M} briefly, let's say that the path $s^i . \alpha_1 . s_1 . \dots . \alpha_n . s_n$ in \mathcal{A} maps to the path $(s^i, 0) . (s_1, 1) . (s_2, 2) . \dots . (s_n, n)$ in \mathcal{M} . For all $0 \leq k < r$, the probability of transitioning from (s, k) to $(s', k+1)$ is $p_{s, w[k+1]}(s')$ (note that these probabilities depend on w). For simplicity, the states (s, r) are stuttering. The initial state of \mathcal{M} is $(s^i, 0)$. The previous definitions for \mathcal{M} imply that the probabilities of the paths in \mathcal{A} having length at most r coincide with the probabilities of the corresponding paths in \mathcal{M} :

$$\begin{aligned} & \Pr_{\mathcal{A}}^w(s^i . \alpha_1 . s_1 . \dots . \alpha_n . s_n^\uparrow) \\ &= p_{s^i, \alpha_1}(s_1) \cdot \prod_{k=1}^{n-1} p_{s_k, \alpha_{k+1}}(s_{k+1}) \\ &= \Pr_{\mathcal{M}}((s^i, 0) . (s_1, 1) . (s_2, 2) . \dots . (s_n, n)). \end{aligned}$$

As a consequence, the probability that w reaches S in at most r steps equals the probability that \mathcal{M} reaches a state in $S \times \{0, \dots, r\}$. The latter probability can be calculated using standard techniques, as it poses a simple reachability problem for finite Markov chains.

We have just showed that l_r is computable. We still need to prove that it complies with the properties we need so that our main algorithm works. In order to prove Inequation 2, we use the fact that $l_r = \min_{w \in W_r} \Pr^w(\mathcal{H}_r)$, where W_r is the set of all words of length r . Let w^* be $\arg \min_{w \in W_{r+1}} \Pr^w(\mathcal{H}_{r+1})$ and w^{*-1} be $w^* \downarrow r$. The required inequality $l_r \leq \Pr^{w^*}(\mathcal{H}_{r+1})$ follows since $l_r = \min_{w \in W_r} \Pr^w(\mathcal{H}_r) \leq \Pr^{w^{*-1}}(\mathcal{H}_r) = \Pr^{w^*}(\mathcal{H}_r) \leq \Pr^{w^*}(\mathcal{H}_{r+1})$, where the last inequality holds since $\mathcal{H}_r \subseteq \mathcal{H}_{r+1}$.

Next, we prove Inequation 3. Let $\mu = (\psi(m))_{m=1}^\infty$ be a sequence of infinite words such that $\lim_{m \rightarrow \infty} \Pr^{\psi(m)}(\mathcal{H}_r) = I$ and the sequence $(\Pr^{\psi(m)}(\mathcal{H}_r))_{m=1}^\infty$ is non-increasing (such a sequence exists by definition of infimum). Let w^* be a word of length r that appears infinitely often in the sequence $(\psi(m) \downarrow r)_{m=1}^\infty$ (this word exists as the sequence is infinite, and there are finitely many words of length r).

We prove Inequation 3 by proving $\Pr^{w^*}(\mathcal{H}_r) \leq I$. Suppose, towards a contradiction, that $\Pr^{w^*}(\mathcal{H}_r) > I$. Then, by definition of μ there exists $\psi(p)$ in μ such that $\Pr^{w^*}(\mathcal{H}_r) > \Pr^{\psi(p)}(\mathcal{H}_r) \geq I$. Since w^* appears infinitely often in $(\psi(m) \downarrow r)_{m=1}^\infty$, there exists $q > p$ such that $\psi(q) \downarrow r = w^*$. Since the values $\psi(m)$ are non-increasing, we reach the following contradiction: $\Pr^{w^*}(\mathcal{H}_r) > \Pr^{\psi(p)}(\mathcal{H}_r) \geq \Pr^{\psi(q)}(\mathcal{H}_r) = \Pr^{w^*}(\mathcal{H}_r)$.

It remains to prove Eq. 4. In order to prove this equality, let Ψ be the sequence

$$(\Psi_r = \arg \min_{w \in W_r} \Pr^w(\mathcal{H}_r))_{r=1}^\infty$$

(the set W_r has been defined above). Note that

$$l_r = \Pr^{\Psi_r}(\mathcal{H}_r). \quad (8)$$

Given Ψ , we construct an infinite *limit* word² $\vec{\psi}$ having the property that, for every M , the prefix $\vec{\psi} \downarrow M$ appears infinitely often in the sequence $(\Psi_r \downarrow M)_{r=M}^\infty$. We take the first symbol $\vec{\psi}[1]$ to be any symbol that appears infinitely often in $(\Psi_k \downarrow 1)_{k=1}^\infty$. In order to obtain the second symbol $\vec{\psi}[2]$, we consider the subsequence Ψ^1 of Ψ containing all words in Ψ whose first symbol is $\vec{\psi}[1]$. Then, $\vec{\psi}[2]$ is any symbol that appears infinitely often as the second symbol in $(\Psi_k^1 \downarrow 2)_{k=2}^\infty$. In general, we can describe the process to obtain Ψ^M and $\vec{\psi}[M]$ in an inductive fashion, by stating that $\vec{\psi}[M]$ is any symbol that appears infinitely often in $(\Psi_k^{M-1} \downarrow M)_{k=M}^\infty$ and Ψ^M is an (infinite) subsequence of Ψ^{M-1} complying with $\Psi_k^M[M] = \vec{\psi}[M]$. The existence of the subsequence Ψ^M ensures that $\vec{\psi} \downarrow M$ appears infinitely often in $(\Psi_r \downarrow M)_{r=M}^\infty$, as desired.

²We use the word *limit* as it resembles the *limit schedulers* in [8].

As an auxiliary result, we prove $\Pr^{\vec{\psi}}(\mathcal{H}) = I$. Suppose, towards a contradiction, that $\Pr^{\vec{\psi}}(\mathcal{H}) > I$. Then, there exists ψ' such that $\Pr^{\vec{\psi}}(\mathcal{H}) > \Pr^{\psi'}(\mathcal{H}) \geq I$. As³

$$\forall \psi : \Pr^\psi(\mathcal{H}) = \lim_{k \rightarrow \infty} \Pr^\psi(\mathcal{H}_k), \quad (9)$$

there exists K such that

$$\Pr^{\vec{\psi}}(\mathcal{H}_K) > \Pr^{\psi'}(\mathcal{H}) \geq \Pr^{\psi'}(\mathcal{H}_M) = \Pr^{\psi' \downarrow M}(\mathcal{H}_M) \quad (10)$$

for all M . By definition of $\vec{\psi}$, there exists $M > K$ such that $\Psi_M \downarrow K = \vec{\psi} \downarrow K$. Then, $\Pr^{\vec{\psi}}(\mathcal{H}_K) = \Pr^{\vec{\psi} \downarrow K}(\mathcal{H}_K) = \Pr^{\Psi_M \downarrow K}(\mathcal{H}_K) \leq \Pr^{\Psi_M}(\mathcal{H}_M) \leq \Pr^{\psi' \downarrow M}(\mathcal{H}_M)$ (where the last inequality holds by definition of Ψ_M) thus contradicting Inequation 10.

Now we are ready to prove $\lim_{r \rightarrow \infty} l_r = I$. Since $l_r \leq I$ for all r , we have $\lim_{r \rightarrow \infty} l_r \leq I$. Suppose, towards a contradiction, that $\lim_{r \rightarrow \infty} l_r < I$. Then, by $\Pr^{\vec{\psi}}(\mathcal{H}) = I$ and Eq. 9, there exists K such that

$$\lim_{r \rightarrow \infty} l_r < \Pr^{\vec{\psi}}(\mathcal{H}_K) = \Pr^{\vec{\psi} \downarrow K}(\mathcal{H}_K). \quad (11)$$

By definition of $\vec{\psi}$, there exists $M > K$ such that $\Psi_M \downarrow K = \vec{\psi} \downarrow K$. Then, by Eq. 8, we have $\lim_{r \rightarrow \infty} l_r \geq \Pr^{\Psi_M}(\mathcal{H}_M) \geq \Pr^{\Psi_M \downarrow K}(\mathcal{H}_K) = \Pr^{\vec{\psi} \downarrow K}(\mathcal{H}_K)$, which contradicts Inequation 11.

2.2. Upper bounds

For our upper bounds, we use *lasso-shaped* words (LSW). A LSW is an infinite word of the form $\psi = \alpha_1 \cdots \alpha_K (\beta_1 \cdots \beta_M)^\omega$, in which the last M in which the sequence of symbols $\beta_1 \cdots \beta_M$ is looped infinitely many times. The name *lasso-shaped* is borrowed from the counterexamples for LTL properties of Büchi automata, this name being used, for instance, in [9]. Such counterexamples also consist of a finite stem and a sequence that is looped infinitely many times.

In this paper, we restrict to LSWs with $M \leq 2^{|S|}$ (recall that S is the set of states of the PFA), and we say that K is the *order* of ψ , denoted by $\text{order}(\psi)$. Note that, because of our restriction on the length of the loop, the amount of LSWs with order at most K is finite.

We denote by $\text{LSW}(r)$ the set of all LSW with order at most r . The set of all infinite words is denoted by Words .

For upper bounds, we take $u_r = \inf_{\psi \in \text{LSW}(r)} \Pr^\psi(\mathcal{H})$. Inequalities 5 and 6 follow from $\text{LSW}(r) \subseteq \text{LSW}(r+1) \subseteq \text{Words}$.

³This equality is standard for reachability properties, and can be deduced from $\Pr^\psi(\mathcal{H}) = \Pr^\psi(\bigcup_{k=1}^\infty \mathcal{H}_k \setminus \mathcal{H}_{k-1}) = \sum_{k=1}^\infty \Pr^\psi(\mathcal{H}_k \setminus \mathcal{H}_{k-1})$.

The computability of u_r follows in a similar way to that of l_r : the amount of LSWs having order at most r is finite, and we can explore the probabilities for each of these words. Similarly as for the lower bounds, the probability for a word $w_1(w_2)^\omega$ is calculated by constructing a finite Markov chain. We just outline the construction. The set of the states of the Markov chain is

$$S \times \{1, \dots, \max\{K, M\}\} \times \{S, L\}$$

(where $K = \text{len}(w_1)$ and $M = \text{len}(w_2)$). The initial state is $(s^i, 1, S)$. In the state (s, n, S) ((s, n, L) , resp.), the probability distribution for the next state is determined by the n -th symbol in the stem (in the loop, resp.). In symbols, the probability of transitioning from (s, n, S) to $(s', n+1, S)$ is $p_{s, w_1[n]}(s')$ whenever $n < K$. From (s, K, S) to (s', K, L) , the probability is $p_{s, w_1[K]}(s')$. The probabilities for the loop are defined in a similar way: the only difference is that in a state (s, M, L) in the end of the loop, we have that $p_{s, w_2[M]}(s')$ is the probability of transitioning to $(s, 1, L)$ (that is, we return to the beginning of the loop). Note that all the paths with positive probability are of the form

$$\begin{aligned} & (s_1, 1, S) \cdots (s_K, K, S) \\ & \cdots (s_{K+1}, 1, L) \cdots (s_{K+M}, M, L) \\ & \cdots (s_{K+M+1}, 1, L), \dots, (s_{K+M+M}, M, L) \cdots \end{aligned}$$

It is easy to see that the probability $\Pr^{w_1(w_2)^\omega}(\mathcal{H})$ is the probability of reaching a state (s, n, l) such that $s \in S$, and so the minimum probability for all words of order at most K can be obtained by constructing a Markov chain for each of such words.

It remains to prove Eq. 7. If $I = 1$, then $u_r = 1$ for all r , and so the equation is trivial. From now on, we concentrate on the case $I < 1$. In order to prove that the limit is the infimum, it suffices to show that, for all ϵ , there exists R such that

$$\inf_{\psi \in \text{LSW}(R)} \Pr^\psi(\mathcal{H}) < I + \epsilon. \quad (12)$$

We can indeed restrict to ϵ such that

$$\epsilon < 1 - I. \quad (13)$$

(Having proved the result for such values, the result also holds for the values ϵ' such that $\epsilon' \geq 1 - I$, by taking ϵ such that $\epsilon = (1 - I)/2 < 1 - I \leq \epsilon'$ and hence $\inf_{\psi \in \text{LSW}(R)} \Pr^\psi(\mathcal{H}) < I + \epsilon < I + \epsilon'$.)

We prove Inequation 12 by showing that there exists $\psi^* = w_1(w_2)^\omega$ with $\text{len}(w_2) \leq 2^S$ such that $\Pr^{\psi^*}(\mathcal{H}) < I + \epsilon$. By taking R to be the order of ψ^* , we obtain Inequation 12, that is, the value u_R is ϵ -close to I .




Figure 1: Avoiding S with high probability

Let $\psi^{\epsilon/2}$ be an infinite word such that $\Pr^{\psi^{\epsilon/2}}(\mathcal{H}) < I + \epsilon/2$ (such a word exists by definition of infimum). Using this word, we construct the word ψ^* with the desired properties. For this construction, we focus on the probability of *not* reaching S (that is, the probability of all infinite paths such that none of the states is in S). By definition of $\psi^{\epsilon/2}$, we know that $\psi^{\epsilon/2}$ does not reach S with probability greater than $1 - \epsilon/2 - I$; in symbols:

$$\Pr^{\psi^{\epsilon/2}}(\neg\mathcal{H}) > 1 - \epsilon/2 - I, \quad (14)$$

where $\neg\mathcal{H}$ is the complement of \mathcal{H} , that is, the set of all infinite paths ρ such that $\rho[k] \notin S$ for all k .

Using $\psi^{\epsilon/2}$, we define ψ^* in such a way that

$$\Pr^{\psi^*}(\neg\mathcal{H}) > 1 - \epsilon - I \quad (15)$$

and so $\Pr^{\psi^*}(\mathcal{H}) < I + \epsilon$. The proof proceeds by finding numbers K and M such that the first $K + M$ symbols of ψ^* are the same as in $\psi^{\epsilon/2}$. We name these symbols $\alpha_1, \alpha_2, \dots, \alpha_K, \beta_1, \beta_2, \dots, \beta_M$. After these symbols, the word ψ^* repeats β_1, \dots, β_M indefinitely. This word is illustrated in Fig. 1. The intuition behind the proof is that there exists a set Q_1 of states such that, after exactly K steps, there is sufficiently high probability to be in Q_1 , without hitting S (in the figure, states in S are represented with crosses). Moreover, if Q_{i+1} (Q_1 , respectively) is the set of all states that can be reached after symbol β_i (β_M , resp.) occurs in some state in Q_i (Q_M , resp.), then $Q_i \cap S = \emptyset$ for all $1 \leq i \leq M$. We find K, M and show that ψ^* complies with Inequation 15.

In order to obtain the required K, M , we profit from the fact that a PFA according to our definition can be seen as a particular case of an MDP. For the sake of completeness, we show how our definition for PFA

matches the definition of MDP in [10]. If the MDP underlying a PFA \mathcal{A} is obvious to the reader, then the rest of this paragraph can be safely skipped. In [10], (Def. 3.1), an MDP $\Pi = (S, A, p)$ is defined by a set of states S , a set of actions $A(s)$ enabled at each state s , and probabilities $p_{st}(a)$ of stepping from s to t using a , for each $a \in A(s)$. When mapping a PFA \mathcal{A} to an MDP Π , the set of states S of Π is the same set of states as in \mathcal{A} ; for each s the set $A(s)$ of actions enabled is the set Σ ; the probabilities $p_{st}(a)$ in [10] are simply $p_{s,a}(t)$.

Using the MDP underlying \mathcal{A} , we can resort to the *end-component theorem* ([10, Thm. 3.2]). In terms of PFA, the definition of an end component is as follows.

Definition 1. An end component is a set $E \subseteq S \times \Sigma$ such that for every states $s_1 \neq s_n$ in (a pair in) E there exists a path $s_1.\alpha_1.s_2.\dots.\alpha_{n-1}.s_n$ such that $(s_k, \alpha_k) \in E$ and $p_{s_k, \alpha_k}(s_{k+1}) > 0$ for all $1 \leq k \leq n-1$. We write $\text{States}(E)$ for the set of states of E . When no confusion arises, we simply write $s \in E$ instead of $s \in \text{States}(E)$.

Let \mathcal{E} be the set of infinite paths $s_1.\alpha_1.s_2.\dots$ such that there exists T such that the set $\{(s_t, \alpha_t) \mid t > T\}$ is an end component. The end-component theorem states that \mathcal{E} has probability 1 for all words. The paths in \mathcal{E} are said to *end* in an end component. Then, the set of paths that do not end in an end component (that is, the paths for which no such T exists) has probability 0 for all words and, roughly speaking, can thus be disregarded in probability calculations.

From now on, we are interested on the set \mathcal{E} comprising all paths ending in an end component. Now we show a partition for \mathcal{E} . For all finite paths π , end components E , let $Z(\pi, E)$ be the set of all infinite paths $\pi.\alpha_1.s_2.\alpha_2.\dots$ such that $(\alpha_k, s_{k+1}) \in E$ for all k . Next, we prove that the set \mathcal{E} is equal to $\mathcal{E}' = \bigcup_{(\pi, E) \in \mathcal{Z}} Z(\pi, E)$ where \mathcal{Z} is the set of all pairs (π, E) such that π is either the trivial path s^i , and $s^i \in E$; or $\pi = s^i.\dots.s_{n-1}.\alpha_n.s_n$ and $(s_{n-1}, \alpha_n) \notin E$ and s_n in E . In words, the last state/symbol pair is not in E , but the last state is. Clearly, the inclusion $\mathcal{E}' \subseteq \mathcal{E}$ holds as the paths in $Z(\pi, E)$ end in E for all π, E . In order to prove the inclusion $\mathcal{E} \subseteq \mathcal{E}'$, we prove that any path $\psi \in \mathcal{E}$ is also in \mathcal{E}' . Since $\psi \in \mathcal{E}$, there exists T as in Def. 1. Let's consider the minimum such T . The existence of T ensures that ρ has a prefix π after which all the pairs state/symbol are in E . Moreover, since we are considering the minimum T , either π is the trivial path s^i , and s^i is in E ; or the last state/symbol pair before $\text{last}(\pi)$ is not in E . In summary, the fact that R is minimum ensures that there exists $(\pi, E) \in \mathcal{Z}$ such that $\rho \in Z(\pi, E)$. It remains to prove disjointness, that is,




Figure 2: (s_1, α) is in E , but (s_2, α) is not

$Z(\pi, E) \cap Z(\pi', E') \neq \emptyset$ imply $(\pi, E) = (\pi', E')$. Suppose that there exists $\rho \in Z(\pi, E) \cap Z(\pi', E')$. The set of all state/symbol pairs that appear infinitely often in ρ are all the pairs in E (as $\rho \in Z(\pi, E)$), and the same goes for E' , thus yielding $E = E'$. It remains to prove $\pi = \pi'$. We have that π and π' are both a prefix of ρ . Moreover, since we consider only finite paths in which the last state/symbol pair is not in E , we have that π is the smallest prefix such that after π all the state/symbol pairs are in E , and the same holds for π' . Then, both π and π' have the same length, and so $\pi = \pi'$.

As a consequence of the partition we found, and the end-component theorem, for all words ψ we have $\Pr^\psi(\Omega) = \Pr^\psi(\mathcal{E}) = \sum_{\pi} \sum_{\{E \mid (\pi, E) \in \mathcal{Z}\}} \Pr^\psi(Z(\pi, E))$. If a path ends in an end component E and does not hit S , then no prefix hits S , and E has no states in S . Hence, for all words ψ we have

$$\Pr^\psi(\neg \mathcal{H}) = \sum_{\{\pi \mid \pi \cap S = \emptyset\}} \sum_{\{E \mid (\pi, E) \in \mathcal{Z} \wedge E \cap S = \emptyset\}} \Pr^\psi(Z(\pi, E)).$$

The outer sum ranges over all finite paths such that no state is in S (which we denote as $\pi \cap S = \emptyset$), and the inner sum ranges over all end components E such that the last state/action pair in π is not in E , the last state is in E , and no state of E is in S (denoted by $E \cap S = \emptyset$). In particular, for the word $\psi^{\epsilon/2}$ in Inequation 14, we have $\Pr^{\psi^{\epsilon/2}}(\neg \mathcal{H}) = \sum_{\pi \cap S = \emptyset} \sum_{(\pi, E) \in \mathcal{Z} \wedge E \cap S = \emptyset} \Pr^{\psi^{\epsilon/2}}(Z(\pi, E)) > 1 - \epsilon/2 - I$. Then, there exists a finite set $\mathcal{B} \subseteq \{\pi \mid \pi \notin S\}$ such that $\sum_{\pi \cap S = \emptyset} \sum_{(\pi, E) \in \mathcal{Z} \wedge E \cap S = \emptyset} \Pr^{\psi^{\epsilon/2}}(Z(\pi, E)) > 1 - \frac{3}{4}\epsilon - I$. Let $B = \max_{\pi \in \mathcal{B}} \text{len}(\pi)$. For the sake of brevity, let \mathcal{V} be the set of all pairs (π, E) such that $\pi \cap S = \emptyset$, and $\text{len}(\pi) \leq B$, and $(\pi, E) \in \mathcal{Z}$, and $E \cap S = \emptyset$, and $\Pr^{\psi^{\epsilon/2}}(Z(\pi, E)) > 0$. Then,

$$\sum_{(\pi, E) \in \mathcal{V}} \Pr^{\psi^{\epsilon/2}}(Z(\pi, E)) > 1 - \frac{3}{4}\epsilon - I. \quad (16)$$

Note that we can restrict to the pairs (π, E) such that $\Pr^{\psi^{\epsilon/2}}(Z(\pi, E)) > 0$, as the pairs with probability 0 do not affect the sum. In addition, by Inequation 13, we have $1 - \epsilon - I > 0$, and so in the sum in Inequation 16 there is at least one positive summand $\Pr^{\psi^{\epsilon/2}}(Z(\pi, E))$.

The desired K, M are now obtained from $\psi^{\epsilon/2}$ and B . Note that, although we restricted to the summands

complying with $\Pr^{\psi^{\epsilon/2}}(Z(\pi, E)) > 0$, it is still possible that \mathcal{A} exits E after π with positive probability (as the same symbol might be inside E for a reachable state s , but outside E for a state s' that is reachable after the same number of steps as s , see Fig. 2). We show that, by considering arbitrarily large paths, the probability that E is exited becomes arbitrarily small.

Let $b_k^w(s)$ be the probability that, after k steps, the state is reached is s . We generalize this notation to sets of states. Formally: $b_k^w(T) = \sum_{\text{len}(\pi)=k, \text{last}(\pi) \in T} \Pr^w(\pi^\dagger)$. We call the distribution $b_k^w(\cdot)$ a *belief state*, following the nomenclature for POMDPs [11]. Since the set of states is finite, there exist two indices $x < y$ such that $\text{supp}(b_x^{\psi^{\epsilon/2}}) = \text{supp}(b_y^{\psi^{\epsilon/2}})$ (where supp denotes the support of the distribution). Moreover, given any two numbers X, Y , such that $Y > X + 2^{|S|}$, we have $X \leq x \leq y \leq Y$ and $\text{supp}(b_x^{\psi^{\epsilon/2}}) = \text{supp}(b_y^{\psi^{\epsilon/2}})$ for some x, y . Since the amount of sequences of the form $T_0\gamma_1 \cdots \gamma_V T_V$ with $V \leq 2^{|S|}$ is finite (where each T_v is a set of states), at least one of such finite sequences appears infinitely many times in the infinite sequence $\text{supp}(b_0^{\psi^{\epsilon/2}})\psi^{\epsilon/2}(1)\text{supp}(b_1^{\psi^{\epsilon/2}})\psi^{\epsilon/2}(2)\cdots$. Suppose this finite sequence is $\sigma = T_0\gamma_1 \cdots \gamma_V T_V$.

We show that we can take $\beta_1, \dots, \beta_M = \gamma_1, \dots, \gamma_V$. In addition, we take K to be a number (defined below) greater than B , in which an occurrence of σ starts.

Given a component E in a pair in \mathcal{V} (defined before Inequation 16), let

$$Q^E = \{s' \in T_0 \cap \text{States}(E) \mid \forall v \leq V : \Pr^{\gamma_1 \cdots \gamma_v}(s' \cdot \gamma_1.s_1 \cdots \gamma_v.s_v) > 0 \implies s_k \in E\}.$$

In other words, Q^E comprises the states in $T_0 \cap \text{States}(E)$ from which, when executing $\gamma_1 \cdots \gamma_V$, we can only reach states in E . Let Q^{-E} be $\text{States}(E) \setminus Q^E$.

Consider the infinite sequence $e(1), e(2), \dots$ of indices where σ starts. We show that $\lim_{v \rightarrow \infty} b_{e(v)}^{\psi^{\epsilon/2}}(s) = 0$ for all $s \in Q^{-E}$. (As the number of states is finite, this implies $\lim_{v \rightarrow \infty} b_{e(v)}^{\psi^{\epsilon/2}}(Q^{-E}) = 0$.) Suppose, towards a contradiction, that for some $s \in Q^{-E}$, $l > 0$, we have $b_{e(v)}^{\psi^{\epsilon/2}}(s) \geq l$ for all v . By definition of Q^{-E} , there exists $s' \notin E$, $d > 0$, and v such that $\Pr^{\gamma_1 \cdots \gamma_v}(s \cdot \gamma_1 \cdots \gamma_v.s') = d$. Then, the probability of staying in E after the n -th repetition of σ is less than or equal to $(1 - (l \cdot d))^n$, for all n . This implies that the probability of staying in E indefinitely is 0, thus contradicting the fact that $\Pr^{\psi^{\epsilon/2}}(Z(\pi, E)) > 0$.

As a result, for all pairs (π, E) in \mathcal{V} , there exists $e(\pi, E) \in \{e(1), e(2), \dots\}$ such that

$$b_{e(\pi, E)}^{\psi^{\epsilon/2}}(Q^{-E}) < \epsilon / (4 \cdot |\mathcal{V}|). \quad (17)$$

Approximate/Compute	Infimum	Supremum [2]
Reachability	$\sqrt{?}$	\times/\times
Acceptance [12]	$?/\times$	\times/\times

Table 1: Existence of algorithms for PFA

Define $K = \max(B, \{e(\pi, E) \mid (\pi, E) \in \mathcal{V}\})$ and $Y(\pi, E) = Z(\pi, E) \setminus \{\rho \mid \rho[K] \in Q^{-E}\}$. We have $Z(\pi, E) \subseteq Y(\pi, E) \cup \{\pi^\dagger \mid \pi[K] \in Q^{-E} \wedge \text{len}(\pi) = K\}$. Then,

$$\Pr^\psi(Z(\pi, E)) \leq \Pr^\psi(Y(\pi, E)) + b_K^\psi(Q^{-E}) \quad (18)$$

for all infinite words ψ .

We have

$$\Pr^{\psi^*}(Y(\pi, E)) \geq \Pr^{\psi^{\epsilon/2}}(Y(\pi, E)) \quad (19)$$

as, under ψ^* , all paths of length K ending in a state in Q^E do not reach states outside E (because of our definition of Q^E and the symbols γ_v). In fact, if $\text{len}(\pi) \geq K$, the scenario in Fig. 2 is possible under $\psi^{\epsilon/2}$, but not possible under ψ^* . Roughly speaking, after K steps the word ψ^* does not escape E , thus yielding higher (or equal) probability for $Y(\pi, E)$ than any word ψ such that $\psi \downarrow K = \psi^* \downarrow K$ and, in particular, than $\psi^{\epsilon/2}$. Then,

$$\begin{aligned} & \Pr^{\psi^*}(\neg \mathcal{H}) \\ & \geq \{Y(\pi, E) \subseteq Z(\pi, E), \text{ the sets } Z(\pi, E) \text{ partition } \neg \mathcal{H}\} \\ & \quad \sum_{(\pi, E) \in \mathcal{V}} \Pr^{\psi^*}(Y(\pi, E)) \\ & \geq \{\text{Ineq. 19}\} \sum_{(\pi, E) \in \mathcal{V}} \Pr^{\psi^{\epsilon/2}}(Y(\pi, E)) \\ & \geq \{\text{Ineq. 18}\} \sum_{(\pi, E) \in \mathcal{V}} \Pr^{\psi^{\epsilon/2}}(Z(\pi, E)) - b_K^{\psi^{\epsilon/2}}(Q^{-E}) \\ & > \{\text{Inequalities 16, 17}\} 1 - \frac{3}{4}\epsilon - I - |\mathcal{V}| \cdot \epsilon / (4 \cdot |\mathcal{V}|) \\ & \geq 1 - \epsilon - I \end{aligned}$$

In conclusion, the word $\psi^* = \alpha_1 \cdots \alpha_K(\beta_1 \cdots \beta_M)^\omega$ (where $\alpha_1 \cdots \alpha_K \beta_1 \cdots \beta_M$ are first $K + M$ symbols in $\psi^{\epsilon/2}$) complies with Inequation 15. Since $\text{order}(\psi^*) = K$, we obtain $\inf_{\psi \in \text{LSW}(K)} \Pr^\psi(\mathcal{H}) < I + \epsilon$. By Inequation 5, this inequality implies $\inf_{\psi \in \text{LSW}(k)} \Pr^\psi(\mathcal{H}) < I + \epsilon$ for all $k \geq K$, thus ensuring Eq. 7.

3. Discussion

Our algorithm is nonprimitive recursive, and we have still nothing to say about the complexity of the problem.

However, the fact that there exists an algorithm to approximate the value is quite surprising considering similar problems for PFA, as shown in Table 1. The table indicates, for the problems of reachability and acceptance, whether there exists an algorithm to approximate and/or to compute extremal values. Note that the only $\sqrt{\cdot}$ in the table corresponds to the result in this paper.

The table also indicates two pending questions: whether there exists an algorithm to effectively compute the infimum for reachability, and whether the infimum for acceptance can be approximated.

The undecidability for the supremum probability has been used to prove that quantitative model checking under partial information [5, 13] is undecidable for properties involving the supremum. The setting of these papers is more general, as several entities might have different information about the state of the system (in contrast, the problem we address in this paper concerns only an environment that has no information about the state of the system). However, we expect that the proof we presented sheds some light on whether this more general problem is computable or not, in case we consider the infimum instead of the supremum.

References

- [1] M. O. Rabin, Probabilistic automata, *Information and Control* 6 (3) (1963) 230–245.
- [2] O. Madani, S. Hanks, A. Condon, On the undecidability of probabilistic planning and related stochastic optimization problems, *Artif. Intell.* 147 (1-2) (2003) 5–34.
- [3] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley, 1994.
- [4] A. Bianco, L. de Alfaro, Model checking of probabalistic and nondeterministic systems, in: P. S. Thiagarajan (Ed.), *FSTTCS*, Vol. 1026 of *Lecture Notes in Computer Science*, Springer, 1995, pp. 499–513.
- [5] S. Giro, P. R. D’Argenio, Quantitative model checking revisited: Neither decidable nor approximable, in: J.-F. Raskin, P. S. Thiagarajan (Eds.), *FORMATS*, Vol. 4763 of *Lecture Notes in Computer Science*, Springer, 2007, pp. 179–194.
- [6] C. Baier, N. Bertrand, M. Größer, On decision problems for probabilistic Büchi automata, in: R. M. Amadio (Ed.), *FoSSaCS*, Vol. 4962 of *Lecture Notes in Computer Science*, Springer, 2008, pp. 287–301.
- [7] R. Segala, Modeling and verification of randomized distributed real-time systems, Ph.D. thesis, Laboratory for Computer Science, MIT, available at profs.sci.univr.it/~segala/www/pub/phd.tar.gz (1995).
- [8] S. Giro, P. R. D’Argenio, On the expressive power of schedulers in distributed probabilistic systems, *Electr. Notes Theor. Comput. Sci.* 253 (3) (2009) 45–71.
- [9] V. Schuppan, A. Biere, Shortest counterexamples for symbolic model checking of ltl with past, in: N. Halbwachs, L. D. Zuck (Eds.), *TACAS*, Vol. 3440 of *Lecture Notes in Computer Science*, Springer, 2005, pp. 493–509.
- [10] L. de Alfaro, Formal verification of probabilistic systems, Ph.D. thesis, Stanford University, available at users.soe.ucsc.edu/~luca/papers/thesis/thesis.pdf (1997).
- [11] L. P. Kaelbling, M. L. Littman, A. R. Cassandra, Planning and acting in partially observable stochastic domains, *Artif. Intell.* 101 (1-2) (1998) 99–134.
- [12] V. D. Blondel, V. Canterini, Undecidable problems for probabilistic automata of fixed dimension, *Theory Comput. Syst.* 36 (3) (2003) 231–245.
- [13] S. Giro, Undecidability results for distributed probabilistic systems, in: M. V. M. Oliveira, J. Woodcock (Eds.), *SBMF*, Vol. 5902 of *Lecture Notes in Computer Science*, Springer, 2009, pp. 220–235.